PROCUREMENT AND MANAGEMENT OF SPARES

H. S. Campbell

July 1966

P-3404

PROCUREMENT AND MANAGEMENT OF SPARES

H. S. Campbell*

The RAND Corporation, Santa Monica, California

---

*Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

This paper was prepared for publication in a monograph, Defense Management, edited by Stephen Enke. The monograph will be published by Prentice-Hall.

Introduction*

It is important to remember that, although the military materiel system is large and diverse, it provides but one of the inputs to the creation of military capability. The fundamental objective of spares management, therefore, is to furnish the desired level of one input to a productive process that also uses maintenance personnel, equipment, and supplies (other than spares). Implicitly or otherwise, some judgment must be formed as to the desired level of the spare-parts input; and it becomes the task of spares management to accomplish this as efficiently as possible.**

This management job is rightfully a major preoccupation of the military departments and the Department of Defense. The Air Force alone stocks about 1,200,000 spare items, with an acquisition value of more than five billion dollars (active inventory, not including spare engines). Additions to spares stocks amount to nearly three-quarters of a billion dollars annually. This inventory creates a managerial task of tremendous scale.

Spares management has many facets -- such as organization, transportation, storage and handling of materiel, repair of parts, and cataloguing -- but the most fundamental area concerns policies for stock control and spares requirements determination. This aspect of spares management will be the main concern of this chapter, with attention focused primarily upon the application of the tools of inventory-control theory to the analysis of spares procurement and management.

---

*Numbers appearing in brackets refer to references listed at the end of the chapter.

**The view taken here is that spares, maintenance personnel and equipment, and the major equipment items themselves (for example, B-52 aircraft) are inputs to a productive process and substitutable for each other within limits. Spares management is then, technically, a problem in suboptimization of the system providing the spares input for any specified level of that input. References [6], [7], and [18] develop these ideas more completely. Various approaches have been suggested or tried for fixing the level of the spares input. These will be discussed in this chapter.

In other words, the discussion will concern the rules for deciding which inventory levels are optimal or preferred at the various stockage points in a system. The focus is dictated by the key role played by stockage decisions in a materiel system. Stockage policy has a pervasive, if implicit, control over the systems operating costs: It determines the level of investment and the frequency of shortages, of requisitions, of new procurements, and of repair actions.

An introductory survey of the subject, such as this, cannot cover the entire field of inventory theory or its applications. Instead, the chapter will be broken down into three parts: (1) a discussion of the main choices in modeling logistics inventory systems and the interpretation of inventory models and their results; (2) a review of some inventory models representing basic analytic approaches that should have had -- or can be expected to have -- the largest role in application to logistics systems; and (3) a presentation of several special topics of importance to the study of spares management.

Inventory Analysis -- Generalities

It is important to realize that inventory models are, as a class, severe abstractions of the real world. As with other decision-oriented models, their value lies in a systematic relating of the main elements of a problem, so that a large volume of recurring decisions of the same type can be made consistently in accordance with agreed-upon criteria. But before an outline of the structure of inventory analyses is presented, some discussion of present-day inventory management may be in order.

Spares management presents a considerable problem of choice to the would-be analyst. He can be concerned with the determination of future requirements or with the distribution of present stocks. He may deal with the entire system of stockage points, including depots or primary supply points, intermediate storage sites, repair points, and final users, or with some subgrouping of points; or he may be able to treat points independently of one another for analytic purposes. In some problems it may be appropriate to deal with spares according to a

property-class grouping (Aircraft Wheel and Brake Systems) and for other purposes identification of the specific weapon (C-130 Aircraft) may be important.

The existing framework for data collection and management control, which must limit an investigation to some extent, will be more or less elaborate, depending on the category of cost and repairability to which a particular set of spares belongs. In the Air Force, for example, although high-cost, recoverable spares represent only some three percent of the number of items, they account for nearly seventy percent of total spares investment. Clearly, more elaborate management methods are justifiable for these articles than for low-cost material. This management framework may be taken as given for many purposes, though ultimately some judgment must be reached as to the appropriate cost of data collection, accounting, and other forms of management input.

As a commodity, spare parts possess several attributes with particular implications for management. Procurement lead-times may be quite long, on the order of a year or more for high-cost spares. Future spares usage is very difficult to predict, partly because of its characteristic randomness, partly because of the need to forecast the operational program of the end-use weapon, and partly because of obsolescence due to design change. Initial estimates are particularly error-prone [19], a fact that is important in studies of the provisioning area.

How much of the total spares-management system must be embraced in a single analytic formulation? It is certainly neither practical nor necessary to deal with everything at once. Simplifications of the problem are obviously desirable from the standpoint of facilitating computation of policies and managerial understanding; however, a valid analysis does require that all interacting features of the problem be included in a meaningful way.

For many purposes, for example, it is important to consider related sets of spares, such as those associated with a particular weapon, rather than to deal with individual items. This broader view is necessary because different spares are, within limits, substitutes for each other as inputs to weapon maintenance. Maintenance efficiency and weapon up-time are sensitive to aggregate supply effectiveness, and single-item models

are not well suited to analysis or display of aggregate costs and outcomes. It is frequently not enough to know that a policy in force possesses optimal properties for each item; it may also be important to determine the relation between total cost and effectiveness for decision-making purposes.

Interaction also exists among echelons. The stock levels held at depot in part determine the depot-base resupply time, which affects the stock levels needed at the base. Further, the frequency of ordering from the base affects the cost of depot operation, through the number of orders the depot must process and the effect on the variance of depot issues. These considerations argue for comprehensive multi-echelon, multibase analyses, but such approaches have met with limited success.

The Structure of Inventory Analyses

The conventional approach is to cast the problem in the form of an economic model, for which either (1) costs are the only relevant consideration and are minimized by the choice of stockage policy or (2) policy is chosen to minimize costs for some given level of effectiveness.* These models must embody assumptions or rules about a number of different components of the real-world counterpart system: a policy structure or format, a particular model of the demand process, a model of the replenishment process, and a statement about the dependence of the relevant costs upon these factors.

The Ordering Policy. The ordering discipline is frequently dictated by the problem or, more rarely, is one of the things to be determined by analysis. By far the commonest is the (s,S), or two-bin, policy, for which a large body of theory also exists. When the stock, x, falls below the predetermined level, s, an order is placed for S-x units, S and s being chosen so as to minimize costs. A related discipline is the (s,Q) policy, in which the order placed when the stock falls to or below s is of a fixed size, Q. These ordering policies may be coupled with a system of continuous review, in which it is known when the reorder point is passed. This implies current knowledge of balances,

*Alternatively, effectiveness may be maximized for some given budget.

which may be a trivial requirement if a single activity is involved but may involve elaborate reporting of transactions in larger systems.

An alternative to continuous review is periodic review. This corresponds to the wholesale-requirements methods in a number of areas where stocks on hand are compared, quarterly perhaps, to desired stocks and the difference placed on procurement each time. Or, a periodic review system may be combined with an (s,S) or (s,Q) policy, in which stock is reordered up to S (or ordered in batches of size Q) if below s at the time of review (otherwise no order is placed).

An important special case of the (s,S) policy is the (S-1,S) policy, in which "one-for-one" ordering is practiced (the reorder point being one less than the maximum stock). When this policy form can be assumed to be optimal (typically, where higher-cost spares are involved), it is only necessary to compute values for a single policy variable, S, and certain other analytic simplifications obtain. This makes possible more comprehensive analysis for this case.

Models of Demand. The representation of the way in which demands for inventory arise is an important part of the analysis of inventory problems. It is usual to treat spares demand as a random process described by an appropriate probability distribution function (inventory studies have also dealt with known, or deterministic, demand, but these are not relevant to the spares problem). The modeling task then becomes one of selecting the best function by considering empirical fit, analytic convenience, and the physical process embodied by the probability model.

The physical assumptions corresponding to the Poisson distribution are suggestive of the conditions under which spares demand occurs. A Poisson distribution arises where: (1) the number of demands occurring in any interval of time is statistically independent of the number occurring in any other nonoverlapping interval (that is, the fact that a demand occurred today gives no information about the number of demands tomorrow), (2) the process is stable over time, and (3) if time is divided into small enough intervals, the probability of two or more demands in the same interval is negligible ([12], pp. 143-154).

The advantages of the Poisson model are compelling: Its assumptions seem to be satisfied in practice; its shape (skewed to the right) corresponds well to empirical frequency distributions of demand; it is analytically tractable; and it is a one-parameter distribution, requiring only the mean rate of demand (or average issue rate) for estimation purposes.

But often the Poisson does not give a good fit to spares-demand data, and some other model of the underlying demand process is chosen [13], [9]. This inadequacy can be attributed to circumstances interfering with the simple Poisson assumptions: There may be "contagion" effects in maintenance if discovery of a defective part on one piece of equipment leads to inspection of other units (and possibly preventive replacement); some parts are liable to damage during installation, and activity of different aircraft is usually correlated so that the exposures to the possibility of failure are bunched in time.

The effect of these, or possibly other, unknown, circumstances is to raise the apparent variance of demand to values inconsistent with the Poisson hypothesis. A two-parameter distribution is then needed that will, in effect, allow a more exaggerated skewness. Compound Poisson distribution functions are a natural choice: They are two-parameter distributions (allowing a fit to any empirically determined mean and variance); they retain some correspondence to the assumptions suggested by the physical process; and they possess some of the analytical advantages of the simple Poisson.* The commoner forms are the geometric, or stuttering, Poisson and the negative binomial. Without discussing the particulars of either distribution, the physical process they describe is that demands arise in bursts, or clusters. The number in any burst is governed by the geometric or logarithm distributions, respectively, and the occurrence of the bursts themselves constitutes a simple Poisson process.

Whatever probability model is used, an implicit problem of estimation exists in the application of the inventory policy. The model may assume that the distribution and its parameter values are known,

*Reference [12], pp. 268-272, contains a description of the compound Poisson family.

but in practice these must be estimated from past data. Estimation errors are one of the most serious causes of poor performance by inventory proposals. It may be advantageous to reflect any uncertainty about the estimate of the mean rate of demand by means of a Bayesian model (for an example, see [10]).* Inventory policies incorporating Bayesian procedures involve considerably increased computation but possess significant advantages. There is no restriction to expressing uncertainty about future demand in a single set of probabilities that may not correspond very well to reasonable assumptions about the demand process; rather, it is possible to view the mean (or other parameters) of the underlying demand distribution as itself subject to uncertainty, expressed in a prior distribution. Demand is then viewed as a two-stage process of a random choice of the parameter(s) of some type of distribution from the prior distribution and a subsequent random draw of the size of demand from the particular distribution selected. The probabilities of any demand level for this two-stage process can be computed, and, when a sample of actual demand data is observed, the prior distribution can be modified by the application of Bayes theorem and the demand probabilities recalculated. This _ex posteriori_ distribution will exhibit a smaller variance, thus reflecting "learning" or reduced uncertainty about the parameter(s) of the basic or underlying demand distribution. Indeed, the major advantage of the Bayesian approach in inventory models is that it provides a systematic way of reflecting increased knowledge into the estimate of the future. Exploitation of this learning in most cases presupposes a multiperiod inventory analysis.

So far demand has been discussed as though it were a stationary process, that is, identically distributed in every time period. Obviously, this assumption does not hold for some of the most important spares-decision situations, because spares usage is related to weapon activity, which changes continuously as weapons phase in and out. It is still possible to assume that the distribution is stationary with respect to some "program element," such as aircraft flying hours, but models reflecting a varying program are necessarily computationally more complex than "steady-state" models.

d

---

*See [23] for a general treatment of the subject of Bayesian statistics.

The case where demand is not stationary, whether due to varying operational programs (per unit time) or other causes, has been treated as a problem in adaptive forecasting through the use of exponential smoothing [2].

The Replenishment Process. The process by which resupply is accomplished must be described by the inventory model. This may involve resupply from a higher echelon, procurement from outside the materiel system, or repair of spares turned in at the time of demand, depending on the echelon of stockage and the category of spare parts.

A considerable simplification in the analysis is possible if the replenishment time can be treated as fixed or known. When it cannot be a known constant, it is represented as a randomly determined delay, in which the response times are independent of each other and of the stock-level position. This delay may be drawn from an empirical distribution, or an analytic function may be fitted to the data. The significance of these alternatives will become clearer in connection with the specific inventory analyses in the next section.

The Relevant Costs. The relevant costs are, of course, those that vary when the inventory policy changes. Since the cost structure permitted by inventory models is rather simple, the estimation of cost parameters requires careful interpretation and judgment. The costs can be classified under several headings: cost of procurement, cost of ordering and shipping inventories to the point of use, cost of "holding," or keeping inventories in stock, and cost of shortages.

Procurement Costs: This is usually taken to be the unit purchase price plus first-destination transportation cost. In some problems price "breaks," as a function of order size, must be considered. In steady-state models, procurement costs or unit costs do not enter directly, though holding cost is usually estimated as a function of procurement cost.

Holding Costs: These are all costs associated with holding stocks in the logistics system: warehousing costs, inventory-taking costs, cost of modification or maintenance performed on stocks in storage, and so forth. These are related to the time the stocks are held and are

typically estimated as a fraction of the value invested in spares. It is also common to include an "obsolescence risk" factor and an interest charge. The interest charge is, in principle, the foregone earnings on the invested sum if it could be devoted to some other, profitable use. Presumably the government's alternative use of funds is debt-retirement, and studies of military inventory systems have generally used the average cost of the public debt as the interest charge. Where it is desirable to reflect a ceiling or constraint on funds, however, the interest charge may be varied to ration the available funds among competing items.

Ordering Costs: These are the costs associated with preparing a purchase order or contract, in the case of procurement of spares, or with preparing a requisition on a wholesale supply point for internal transfers. They are incurred once per purchase or order. Average cost of shipping materiel from depot to base is sometimes treated as part of base ordering cost.

The Shortage Cost: A distinction is made here between the shortage cost and the shortage penalty. Later in this section, the penalty will be seen to be equivalent to a policy variable or control; however, there are also objectively measurable costs associated with unfilled demands. These are the extra costs of rapid or premium ordering, transportation, or repair actions that are taken in response to a shortage.

A few words on cost estimation. The overall costs of the inventory activity depend on such things as the volume and value of the stock warehoused and the frequency of orders. If these costs are estimated by standard cost-accounting methods and used in single-item inventory formulas, the overall level of activity is in turn determined, and the cost estimates may change. This circularity can lead to adjustment effects when inventory decision rules are installed.

This is one of the more persuasive arguments for dealing with systems or sets of items rather than with individual items. Management is enabled to focus on overall costs and outcomes for some large part of the supply activity, reaching the desired position directly rather than by trial and error. (For a discussion of this point, see [8], p. 76.)

Other types of costs may be important, especially if changes to the spares-management system itself are being studied. Chief among these are the costs of data processing accompanying more elaborate control systems (central accounting for assets, transactions reporting, or demand forecasting techniques). Also important may be the increase or decrease of costs accompanying changes in response times -- requisition processing, transportation, or repair times.

Models and System Objectives. The point has already been made that spares supply is a technical input to a productive process and that the managerial problem has two aspects: to determine the desired level of that input and to provide that chosen level in the more efficient way. The study of logistics has been much more fruitful in the latter task than in the former: Determination of the preferred level of spares input is not well understood as an analytic process (although the problem can certainly be stated formally).

There are essentially two ways of linking the spares subsystem to the overall logistics system. One is to include in the cost function a depletion penalty or shortage penalty representing the cost of compensating for the shortage, over the long run, with increased amounts of other inputs (this procedure will be described more precisely below). Minimization of this cost function with respect to the policy variable will, under the assumptions of the model, provide a stockage policy that is consistent with the desired overall results.

The other method is to require the inventory policy to achieve some specified rate of depletions or shortage, by item or overall. This approach is less commonly used but has several advantages. The rate of shortage occurrences (or their duration, or some similar index of merit) can be compared directly to the occurrence rate in the live inventory system, which allows the performance of the policy and the system to be monitored and permits managerial experience to be applied, directly and naturally, to choosing an appropriate level of effectiveness.

Fundamentally, choosing the effectiveness level and estimating the shortage penalty involve very much the same considerations and information. Consider the following simplified two-part inventory system. The

costs of each part, C_1 and C_2, and the number of shortages of each for some time period, N_1 and N_2, are both functions of the stocks, q_1 and q_2.

Since it is somewhat more convenient to work in terms of the expected number of demands satisfied than expected shortages, we define:

(1) $$Z = \bar{X} - N$$

where $\bar{X}$ is the total expected demand and N the total expected shortage. The minimization of the spares budget, B, can be displayed graphically as in Fig. 1. The curves Z_1,Z_2,/and Z_3 represent several of the many possible isoquants showing the combinations of q_1 and q_2 that produce the given value of Z. The line BB is the combinations of q_1 and q_2 that can be purchased with a given budget (costs are here assumed to be linear with the q's). The largest value of Z that can be obtained

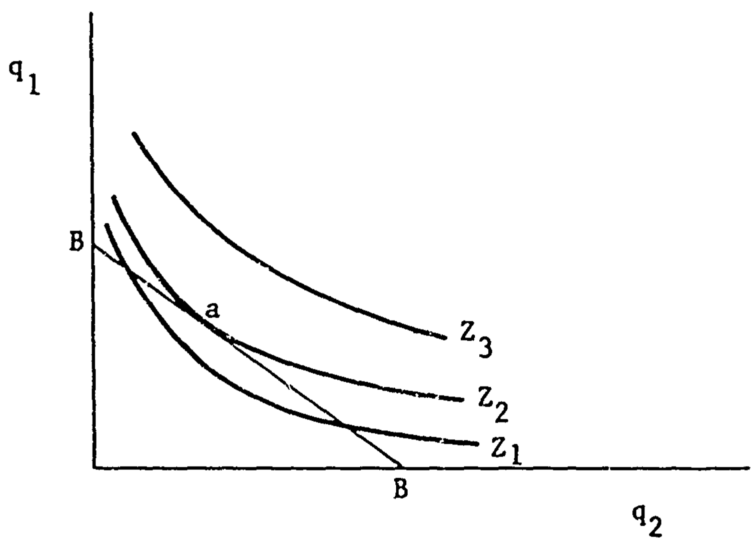

Fig. 1. Minimization of Spares Budget for Given Expected Shortage Rates

with a given budget, or the lowest budget that will achieve a particular level, is obviously at a point of tangency of a budget line with a Z isoquant. This occurs in Fig. 1 at the point a, which is the lowest budget line that will permit the level Z_2 to be reached. At this point of tangency, the rate at which q_1 is being exchanged for q_2

along the constant-Z curve, Z_2, is equal to the rate of exchange of q_1 for q_2 along the constant budget line. Mathematically, this is equivalent to the following:

(2) $$\frac{\frac{\partial Z}{\partial q_2}}{\frac{\partial Z}{\partial q_1}} = \frac{\frac{\partial B}{\partial q_2}}{\frac{\partial B}{\partial q_1}}$$

(3) $$\frac{\frac{\partial B}{\partial q_1}}{\frac{\partial Z}{\partial q_1}} = \frac{\frac{\partial B}{\partial q_2}}{\frac{\partial Z}{\partial q_2}} = \lambda$$

Equation (3) simply states the familiar condition for minimum cost of a given product: marginal cost of inputs should be proportional to marginal physical products.

If λ is now used as a shortage penalty, it is possible to write a conventional cost function for either of the two spare parts.

(4) $$C(q) = B(q) + \lambda\, N(q)$$

The choice of q will minimize C(q) by taking the derivative and setting it equal to zero.

(5) $$\frac{dC}{dq} = \frac{dB}{dq} + \lambda\frac{dN}{dq} = 0$$

Since $\frac{dN}{dq} = -\frac{dZ}{dq}$, Equation (5) is easily seen to be identical with (3). Thus it is possible to use λ as a shortage cost in the set of single-part equations (4) or assign a value of N to each part as in (3). Either procedure yields the same stockage policy, and both require the same considerations, information, and analysis.

Of course, the aggregate shortage objective, N (or Z), must be arrived at in some fashion. By an extension of the foregoing analysis, it is possible to show, conceptually, how it derives from the relationship of inventory inputs to other logistics inputs.

Now consider some overall output, F (which might be in-commission weapons, flying hours, or cargo ton-miles). For illustrative purposes, F will be determined by two input factors, the inventory-system input, Z, and one other, M, which represents perhaps balanced doses of maintenance manpower and equipment. Production isoquants and constant outlay curves will again be used.

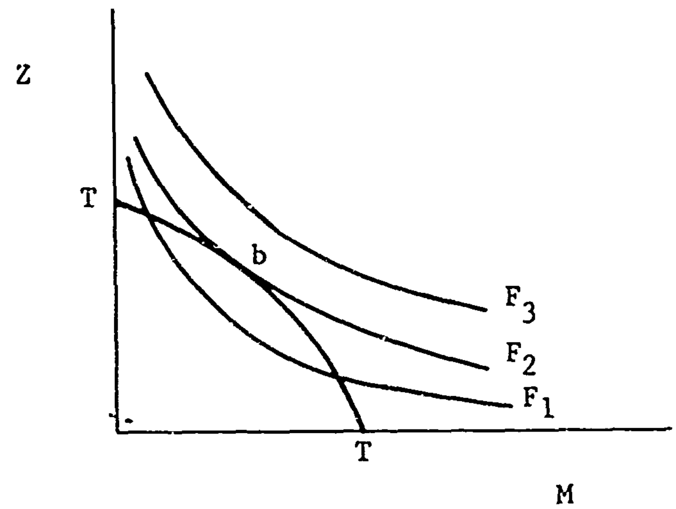

Fig. 2. Minimization of Overall Logistics Budget for Given Levels of Output

By the same reasoning as before, the minimum-cost point is at b for output F_2 (or alternatively, F_2 is the maximum output for the budget represented by TT), and it is characterized by these relations, where T is the total cost.

$$\frac{\frac{\partial F}{\partial Z}}{\frac{\partial F}{\partial M}} = \frac{\frac{\partial T}{\partial Z}}{\frac{\partial T}{\partial M}} \tag{6}$$

$$\frac{\frac{\partial F}{\partial Z}}{\frac{\partial T}{\partial Z}} = \frac{\frac{\partial F}{\partial M}}{\frac{\partial T}{\partial M}} \tag{7}$$

Again, the interpretation of (7) is that all marginal physical products should be proportional to marginal costs, as a condition of efficient production. Since $\frac{\partial T}{\partial Z}$ represents the rate of increase of

total cost as the expected number of demands met increases, this value is simply the imputed value of satisfying a demand under equilibrium conditions. Thus it is the negative of the shortage penalty, which has been designated λ. With the substitution, the result is:

$$\frac{\frac{\partial F}{\partial Z}}{-\lambda} = \frac{\frac{\partial F}{\partial M}}{\frac{\partial T}{\partial M}} \tag{8}$$

$$\lambda = -\frac{\partial F}{\partial Z} \frac{\frac{\partial T}{\partial M}}{\frac{\partial F}{\partial M}} \tag{9}$$

Equation (9) allows an interpretation of the shortage penalty, λ, in the context of overall resource allocation. It is the value of a small change in the overall product due to a change in Z (that is, $\partial F/\partial Z$), measured by the marginal cost per unit of product in terms of other resource inputs -- in this case "maintenance and equipment." The ratio $\frac{\partial T}{\partial M}/\frac{\partial F}{\partial M}$ is the rate of increase in total cost per unit added of the M resource, divided by the rate of increase in the product, F, per unit of M. This is the marginal cost per unit of F. Thus the shortage "cost," λ, can be derived from the optimizing solution for the overall logistics system.

This illustration has been carried out in terms of two parts and two classes of input for illustrative purposes, but in principle it can be generalized to any number of parts and inputs.* Note that λ does not in theory involve the valuation of military worth or the utility lost due to weapon down time; it is determined by the marginal relations between inputs characterizing the minimum-cost (or maximum-output) solution of the logistics resource allocation problem. One has the choice of using the determined value of N (or Z) or the ratio, λ. Examination of (9) shows what is involved in treating λ as a cost

*It has been assumed that the production relation and the constraint are concave functions that are continuous and differentiable and that all inputs are used in positive amounts. A deeper analysis must allow for departures from these assumptions.

to be estimated by engineering and accounting means. One must be able to assume that the entire system has been approximately optimized, to pick one of the other resource inputs, and to evaluate the derivatives involved: the rate of increase of output per unit of the resource added, rate of increase in cost per unit of this resource added, and the rate of decrease of output per unit shortage increase. In general, it is not correct to use the average cost of some input per unit output; for example, the cost of a missile divided by its average alert hours, as a shortage penalty (although this may be an lower bound on the true value of $\frac{\partial T}{\partial M}/\frac{\partial F}{\partial M}$).

Bear in mind that, in practice, the general optimization of a logistics budget is rarely attempted. Efficient use of resources comes about by long-run adaptation, through trial and error. It should be the aim of a spares-management policy to provide "acceptable" service in the most efficient way and, if possible, to present data on the relevant range of efficient alternatives. This will frequently mean that some measure of physical results -- back-order rates or number of shortages -- is more meaningful than an estimate of the ratio, λ, itself.

Inventory Models

Inventory models can be roughly classified by two principles: They may be steady-state or dynamic, according to their assumptions about the future, and single-point or multipoint (including multi-echelon). Another important distinction may be made between single-item and multi-item models. Most applications to logistics problems have been based on single-item, single-point, steady-state models.

As was emphasized in the previous section, the underlying problem in designing a spares-management system is to choose an appropriate activity or level for optimization and to bring into the solution a means of relating stockage policy to an effectiveness objective. We will now examine an important policy area and an inventory model for this area.

The (S-1,S) Inventory Policy in a Multi-item System

The (S-1,S) inventory policy with continuous review lends itself to an accurate analytical formulation and the multi-item approach, and it represents an important policy in logistics applications. High-cost, repairable-type spares are stocked according to this policy, at least at base or final-user level, and represent, as we have observed, a dominating fraction of total spares investment. One-for-one ordering is the desirable policy when the part in question is so expensive that the economic lot size is one unit.

In the repairable-item problem, replenishment of inventories may be either by local repair in an average fraction of the cases, α, or by resupply from a higher echelon in the remaining fraction, $1 - \alpha$. Both repair and resupply times are characteristically random. Since order size is not a variable, the relevant costs (that is, the cnes affected by the choice of S) are the shortage penalty and the cost of holding; the relevant policy outcomes are the number of shortages and the total investment.

A sensible policy would be one that provided less stock for items with short average-response times (perhaps because α was near one) than

those with long times, less for very expensive items than cheaper ones, and less for low demand items than active ones.

The solution of this problem requires calculation of the probability distribution of the number of items in "resupply": that is, the number that have been placed in repair or on order. If this number goes above the spare stock, S, one or more shortages occur. In the most general form of such a problem, the probability distribution of the number of units in resupply must be estimated by Monte Carlo methods or approximated numerically, involving a large amount of computation; however, if the demand-probability function is restricted to the Poisson family (simple or compound), it has been shown that the distribution of the number of units in resupply is of the same Poisson type as the demand and depends only on the average resupply time [9]. Thus any arbitrary (random) resupply distribution can be analyzed quite simply for this class of demand distribution.

This permits a rather extensive treatment of the (S-1,S) policy for a multi-item system, under steady-state conditions and compound Poisson demand [9], [10], [11]. One can formulate quite simply a number of alternative performance measures for an item; one of these is the expected number of shortages* at any point in time, N, as a function of the stock, S.

$$N(S) = \sum_{x=S+1}^{\infty} (x-S)\ P(x;R) \tag{10}$$

P(x;t) is the compound Poisson probability of observing a demand

---

*The referenced studies consider three measures: "fills," the expected number of demands per time period that can be filled immediately; "units in service," the expected number of units in routine resupply at a random point in time; and "ready rate," the probability that the item, observed at a random point in time, has no back orders. "Expected shortages" is the complement of "units of service," which is used here because the earlier discussion of inventory performance has been in terms of shortages rather than fills or units in service.

of x in an interval of time t, and R is the average resupply time, defined above.

Now let us consider a related set of high-cost, repairable items, perhaps those applicable to a particular weapon system. The expected cost is given by the following, where c is the unit cost, h the annual holding-cost rate, and d the shortage penalty.

$$C = dN(S) + hcS \tag{11}$$

Shortages are assumed to be back ordered (this chapter will omit "lost sales"). The minimum-cost value of S is found by taking the derivative of C with respect to S and equating it to zero (or more properly, the difference of C, since S is a discrete variable). This provides the following approximate formula for minimum cost which implicitly determines S.

$$-d\left[N(S+1) - N(S)\right] = hc \tag{12}$$

This says that we should increase S until the value of the decrease in expected shortages just equals the increase in holding cost of adding one more unit to the stock. Since this equation never holds (because S is discrete), the exact procedure is to locate S and S+1 by the following and evaluate C for both.

$$\left[N(S-1) - N(S)\right] > \frac{h}{d} \geq \left[N(S) - N(S+1)\right] \tag{13}$$

The inventory model involved here is certainly easy to understand and use, and it is worthwhile to stop and briefly summarize the assumptions involved. Besides the form of the policy and the particular demand-probability distribution, the model allows for no interactions between stock levels and repair time: The value of S given by (12) does not anticipate use of priorities in out-of-stock situations. Also, the possibility of parts queueing in repair is neglected -- repair times are independent random draws. These are typically the

kinds of assumptions accepted by inventory models in order to gain the advantages of a systematic, computable stockage-decision rule.

There are, however, more important objections to this single-item formulation, for it assumes that h and d are perfectly known and that minimizing these costs for each item satisfies the objectives for the whole set. As the discussion of cost estimating suggested, the "physical" holding-cost portion of h is apt to be affected by the overall storage requirements of the system, and the interest-charge portion depends on the availability of funds. The shortage penalty, d, is more usefully viewed as a device for regulating the number of shortages than as an objectively measured cost. Finally, even if sensible values of h and d are presumed, the solution of Equation (12) for any (or every) item reveals nothing directly about the aggregate outcomes for the set of items. Therefore, a further step is indicated.* Note that the stock level for any item depends on the ratio h/d, which is designated K. If K is treated as a policy-control ratio and applied to all items, there is provided an efficient stockage policy for any values of h and d that are consistent with the ratio. By varying K over an appropriate range it is possible to generate a family of efficient policies and to display the aggregate results in the fashion of Figure 3. This shows the expected number of shortages as a function of investment for the system of related items, where each point on the curve corresponds to the set of policies generated by a particular choice of K.** The expected shortages approach zero asymptotically. Each set of policies generated by a particular K is

---

*For the sake of brevity, there is no discussion at this point of an important aspect of the studies referenced in this section and inventory analysis generally: The problem that the parameters of the demand distribution are not known with certainty. The particular work summarized here included an extensive Bayesian treatment of the matter of demand uncertainty [10].

**The mathematical formulation and computer procedures for generating the relation can be found in [11]. This paper deals with effectiveness measures that approach 100 percent with increasing total investment, but the two procedures are quite analogous.

efficient, for the model chosen: Every item in the system is optimized for that value of the control ratio, and no reallocation of investment among the items can reduce the number of shortages.

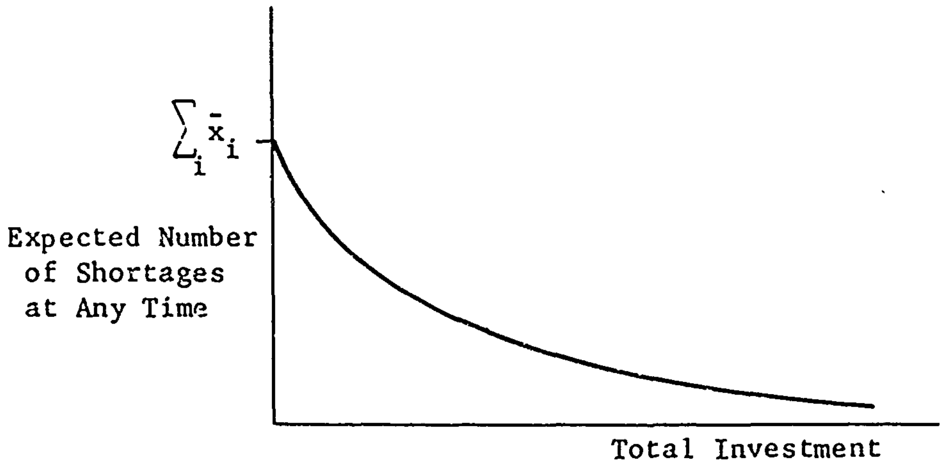

Fig. 3. Shortages vs. Investment for a Set of High-Cost Parts

A manager can select h and d and see, in Fig. 1, the broad consequences of his selection. Alternatively, and more directly, he can simply select the point on the curve of relationship that represents either the available funds or the highest acceptable rate of shortage. He implements this policy with the implicit value of K, which becomes a pure control parameter.

In application, a particular set of items must be chosen to which the computation is to apply. Since it is a single-point formulation, each base, ship, or other using activity must be treated separately (though we may wish to aggregate the results). It may be appropriate to deal with all the (S-1,S) type items at, for example, an air base, as one set; or it may be better management to deal with each major weapon system or mission separately.

In this connection, it might be noted that one of the simplifications of the model as presented here is that all shortages have been treated as being of equal importance, so that their aggregate number has a definite meaning to management. This would not be true

if different parts have different mission essentialities (although high-cost parts are typically important to the mission), in which case essentiality weights might be assigned to each part and the weighted shortage function minimized for each investment level. This is mathematically simple to reflect in models of this type, but deriving essentiality weights is a major undertaking.*

It is probably best to treat major differences in importance associated with different missions, or weapon as opposed to nonweapon uses, by defining different systems or sets of items accordingly.

The Economic Order Size, Reorder Point Model

The steady-state, (s,Q) model of a single item at a single point is the most frequently applied inventory analysis in the logistical area, and elsewhere. The particular formulation shown here is based on [13], but very similar versions can be found in [26] and elsewhere.

In contrast to the (S-1,S) policy, (s,Q) policies are suitable for low-cost spares, in which reordering by batches is desirable, and where, by reason of the large number of items and their low unit cost, it is good management strategy to limit the amount of individual attention and exception actions.

Now, consider the cost relevant to policy determination where an amount Q is to be ordered whenever the stock level falls to (or below) s. Reorders cost r each and are received after a random delay p, subject to a probability distribution h(p). N_s is the expected number of shortages associated with any reorder cycle if the reorder point is s. Other expectations are denoted with a bar over the variable, and the annual cost is written as a function of s and Q as follows ($\bar{x}$ is the average annual demand).

---

*Similarly, we have assumed that additional shortages for the same part have the same importance. Very little has been done with nonlinear shortage penalties in inventory analyses.

$$C = r\frac{\bar{x}}{Q} + hc\bar{y}(s,Q) + dN_s\frac{\bar{x}}{Q} \tag{13}$$

$\bar{x}/_Q$ is simply the number of orders per unit time. An approximation of $\bar{y}$ is $Q/2 + 1/2 + s - \bar{x}_p$, if $\bar{x}_p$ is the expected demand during a reorder time. This neglects the possibility of periods of zero stock in which the system is experiencing back orders. Whether or not $\bar{x}_p$ is subtracted depends on when costs begin to accrue, upon receipt or placement of the order. A more refined model might charge costs against stock in pipeline at a different rate than stock on hand. Equation (13) can be written more explicitly as follows:

$$C = \frac{\bar{x}}{Q}(r + dN_s) + hc(\frac{Q+1}{2} + s - \bar{x}_p) \tag{14}$$

Calculation of N_s requires the marginal distribution of demand during the randomly distributed reorder time. If $g(x)$ is this distribution and $P(x;t)$ is the probability distribution of demand in any time of length t, then the marginal distribution is the following:

$$g(x) = \int_{p=o}^{\infty} P(x;p)h(p)dp \tag{15}$$

The distributions, $P(x;p)$ and $h(p)$ can be of any form, but the computation of probabilities is simplified if these are chosen so that $g(x)$ is analytically specified. For example, if $P(x;t)$ is Poisson, and $h(p)$ can be represented by the gamma distribution, $g(x)$ will be negative binomial in form. N_s is given by (16):

$$N_s = \sum_{x=s+1}^{\infty} (x-s)\ g(x) \tag{16}$$

This formulation assumes that the spacing between orders is large enough so that orders always arrive in the same sequence in which they were placed or that crossing of orders is permissible.

The objective of policy is to minimize (14) by the choice of s and Q, which is done by finding the first differences of C with respect to s and Q and setting them equal to zero. These can be rearranged to yield the following formulas: s is implicitly determined by (18), and Q and s are solved iteratively, first finding Q on the assumption that dN_s is very small and then finding s using this Q, resolving for Q

(17) $$Q = \sqrt{\frac{2\bar{x}(r+dN_s)}{hc}}$$

(18) $$N_{s-1} - N_s > \frac{hc}{d\frac{\bar{x}}{Q}} \geq N_s - N_{s+1}$$

with the new s, and so forth, until convergence is reached. Iteration is not necessary if dN_s is small enough relative to r, in which case the order quantity is determined independently of the reorder point.

These formulas are easily interpreted. Q, obviously, should be inversely related to holding cost, but directly related to the costs associated with reordering (since the larger Q is, the less frequently these costs will be incurred). The logic of the formula determining s is interpreted in the same way as the formula for S in the (S-1,S) case, except that the shortages are incurred only periodically (every $\bar{x}/Q$ years). If (18) is rewritten as $d\frac{\bar{x}}{Q}(N_s - N_{s+1}) = hc$, it can be seen that the object is to balance an annual decrease in cost of shortage against the annual holding cost of one unit.

This type of model is applied to the various stockage points of large systems, taking advantage of the relative independence of these points. Q and s should be large enough, however, to make shortages rare and keep the possibility small that a depot-level shortage will react on the lower level with a series of unusually long resupply times.

The model has a number of assumptions in common with the (S-1,S) model: It assumes that resupply times are random, that the probability distribution of demand is known, and that the present demand conditions

prevail for the indefinite future (the steady-state assumption). These last two assumptions are always violated in practice to a greater or lesser extent: In particular, the steady-state assumption may be quite significant for the depot application because the order quantity may well cover a period of years. During this time bases may enter or leave the program as far as particular items are concerned, and the average demand level will vary accordingly. Base programs are usually level for a number of years, and the Q value will typically be much less than at the higher level ($\bar{x}$ and r are smaller). The life of the item will be longer at depot level, however, and applications of the model can utilize rules for adjusting the order quantity as the end of the program approaches.

The (r,Q) model can be more elaborately and precisely stated than in the foregoing version. A true (s,S) policy computation might be preferable: It is optimal for this problem and often will be the ordering format in use. The (r,Q) model ignores the "overshoot" of the reorder point when more than one demand occurs at the moment the reorder point is reached. In general, however, the system costs are not very sensitive to small errors in s and Q, and there is less at stake in the low-cost area than with the high-cost parts to which the (S-1,S) analysis applies. The limitations of lot-size reorder-point formulas have more to do with their scope than their accuracy as such.

The same objections or difficulties to an item-by-item approach, related in the discussion of the (S-1,S) analysis, apply to the economic order-quantity area as well. Minimizing each item's cost function may fail to minimize total-system costs because the cost parameters are not really constants when total activity varies, and the item-by-item approach tells management nothing about the important aggregate outcomes. In the low-cost area, there may be, in the short run, constraints on manpower or machine capacity for procurement actions or requisitions as well as constraints on funds.

One approach is to extend the control-parameters approach exploited in the (S-1,S) multi-item model. If dN_s is neglected as being probably small where economic order-quantity materiel is

involved (this, at any rate, is a common simplification of the problem), (17) and (18) can be rewritten as (19) and (20) below:

$$(19)\qquad Q = K_1\sqrt{\frac{\bar{x}}{c}} \qquad\qquad \text{where } K_1 = \sqrt{\frac{2r}{h}}$$

$$(20)\qquad N_{s-1} - N_s > \frac{K_2 c}{\frac{\bar{x}}{Q}} \geq N_s - N_{s+1} \qquad \text{where } K_2 = \frac{h}{d}$$

K_1 and K_2 may be interpreted as cost ratios for understanding the general logic of the system but used as controls to determine the overall outcomes for the set of items. K_1 then fixes the ordering frequency and the average investment in order-quantity stocks, and K_2 fixes the expected shortages (given the ordering frequency) and the investment in lead time plus safety-level stocks. The resulting stock levels are efficient in that they optimize stockage for this model and for the particular ratios of costs in the corresponding values of K_1 and K_2 (recall, however, that this simplified formulation is a further approximation of the model of (13); a more complete scheme would allow for the interaction between s and Q).

Management may simply rely on observing the effects of its choices of K_1 and K_2 from accounting reports of the supply activity and make adjustments of the ratios to achieve desired long-run results. Additional apparatus is required if management is to be provided with aggregate data for decision-making purposes. This might involve display of overall outcomes after the fashion of Fig. 3, except that three variables are involved: Number or rate of shortages, number of reorders, and investment. Again, management might use this knowledge of the system's behavior in two ways: To check the results of applying the "best estimates" of the cost factors or to relate item policy systematically to his overall constraints and objectives.* Such systems for providing management with overall information on the outcomes resulting from the choice of stockage policy have not been widely used because of the computational work load involved.

---

*An application and discussion of this general approach is found in [8]. See also [4].

Dynamic Models

Steady-state models are useful for many situations in which the time horizon, or program end, is in the future and demand-rate changes are slow enough for the policy to adapt to them. The limitations of the steady-state approach may be overcome with dynamic programming models, although these are of a higher order of computational complexity. Such models are of great theoretical importance [1], although their application in logistics lies largely in the future.

The basic approach is to divide time into periods that are independent of one another with respect to demand. Demand need not have the same probability distribution in each period, and any desired program length and activity profile can be directly incorporated. The basic analysis involves defining costs in any one period as a function of demand, stock on hand, and amount ordered. The procedure is to solve the model for the final (n^{th}) period, obtaining expected cost as a function of the inherited stock (and the final order in the case of zero lag in delivery of orders). The cost and ordering rules of period n-1 can then be determined, which include the (discounted) costs associated with any level of stock carried forward to period n. The computation continues recursively until the expected cost and ordering rule for the initial period are obtained.

The dynamic formulation is necessary to obtain exact (s,S) policies for the case of periodic review. It is also the basis for the proof that the (s,S) form is optimal under fairly broad conditions, namely that the expected cost of carrying inventory and backorder incidence in each period is a convex function of the starting stock. Unit cost must be constant (in any period), and a fixed cost of placing an order is permitted [24].

Dynamic programming models are very flexible, and can accommodate relatively elaborate systems. Increasing the complexity of the system considered makes it more difficult to generalize the solutions, however, and computation of specific solutions may become costly, even with large-frame computers.

An example of a large, dymamic model will be given in the following section on the multiechelon problem.

The Multiechelon Problem

In view of the circumstance that logistics-supply systems are multiechelon systems,* it is perhaps sobering to note that very little has been accomplished with the analysis of such systems (although they are currently the object of study in several quarters). Interactions among echelons and between repair scheduling and stockage performance can certainly be uncovered, however, that attest to the practical importance of developing this area.

One type of interaction arises because the stock level needed at a lower echelon to provide a stated protection against shortage is a function of the response time (or resupply time) from the next higher echelon, and this response time is in turn a function of the stock level at the higher echelon. Repair scheduling at depot similarly interacts with stockage requirements. Another interaction arises from the determination of base-level order sizes independently of the higher echelon. The frequency of ordering directly affects higher-echelon operation costs, and the size of the order affects the apparent variance of depot issues and increases the stocks required accordingly.

One n-period dynamic, multiechelon model has been formulated, however, the optimality of its policy rules investigated, and some solutions computed [3], [17]. This analysis represents perhaps the most ambitious attempt to date to incorporate a substantial system of higher- and lower-level installations -- depots and bases -- in a single cost-minimizing computation. The model generally applies to items for which there is no ordering or setup cost for shipment within the system, that is, high-cost items for which one-for-one

---

*There are exceptions to this statement. From the standpoint of the Defense Supply Agency or any of the military departments, they are single-echelon systems for DSA items. From the standpoint of the Department of Defense as a whole, of course, these organizations comprise a two-echelon system.

ordering is the accepted discipline. It does permit a fixed cost independent of order size at the highest echelon, however. Procurement and shipping costs, holding costs, and shortage costs are assumed to be linear. Demand originates at the lowest installation or level, and shortages at each echelon are back-ordered.

The policy computation provides the levels to which it pays to order for each point in the system, for each time period in the item's life. The conditions under which Bayesian estimation can be employed with this model have also been investigated [14].

The Clark multiechelon model demonstrates that optimal policies, or, at least, policies that are preferred within some reasonable class, can be computed for extremely complex systems of stockage points. Such techniques involve a considerable investment in data assembly and computer time; and one question involved in considering an application is whether the future can actually be specified in the great detail required with enough certainty to justify the effort.

It might also be noted that the technique just briefly summarized relates to a single item. Little consideration has been given to imposing multi-item constraints on such relatively elaborate one-item models. However, the static multi-item analysis for the (S-1,S) ordering policy under conditions of compound Poisson demand, described earlier in this section, has also been extended to the multibase, multiechelon problem [25]. In the larger context, the problem becomes one of allocating successive units of each spare item to one of the bases, or to the depot, so as to achieve the greatest marginal increase in performance. Performance is measured at the lowest echelon, or base level, by an aggregative criterion such as fill rate or the expected number of outstanding back orders. Back orders or non-fills at the depot echelon are not counted in the criterion, but the allocation of units of stock reflects the fact that a unit assigned to the depot does affect base-level performance by way of the average depot-to-base response time. Allocation of an investment budget among the spare items is then accomplished by marginal analysis, exactly as in the simpler, single-point model.

Special Topics

Deferred Procurement and Phased Provisioning

Perhaps the most vulnerable aspect of the application of decision models to inventory problems relates to the knowledge that must be assumed about demand. In practice, demand probabilities are not known quantities but must be estimated, and this becomes an important source of error in calculating inventory policies. Indeed, much careful optimizing of stockage decisions is simply wasted because of the poor precision with which the parameters and form of the demand distribution can be determined.

This is a particularly important problem in the case of initial stockage decisions for the new spares associated with a weapon-system phase-in. These decisions, which frequently represent a large part of the cotal spares buy, must be based on technical or engineering estimates, made prior to any actual operational experience with the weapon. Such estimates are usually conservative -- that is, they overstate the demand activity that may be expected -- and are at best subject to considerable error [19]. The conservatism, or bias, is apparently a product of the difficulty inherent in making such estimates and of the realization that large underestimates on one or a few important parts could cause serious problems for the new weapon system.

One line of attack on this problem is to attempt to refine the methods by which such estimates are formed or to devise ways of using the information that they do contain. A different approach is represented by the set of ideas or concepts called "deferred procurement" [21], "responsive production," or most recently, "phased provisioning." These systems differ in detail but have a common principle: The substitution of increased management inputs, information processing, and contractor responsiveness for some portion of the spares that would ordinarily have been bought. Under phased provisioning, the current system, the manufacturer advances the production date on units being assembled for eventual incorporation in the end item. This provides a production-line "float" while the weapon concerned

is in production. The support manager then buys only a portion of his total computed requirement, drawing on the float when -- and if -- the need materializes.

Such arrangements do have a cost to the user and quite possibly would not represent the most efficient way of doing business if the demand rate were more certainly known. But under the conditions of relative uncertainty in a weapon's early life, they represent a way of buying some time to form better demand estimates before making a commitment on the entire spares buy. Thus, they may contribute materially to better system stockage decisions.

Systems Analyses

So far the discussion has mainly concerned the methods and choices involved in analyzing the problem of stocking a given support system in as efficient a manner as possible; however, an immediate object of a study of spares management may be to determine the desired performance characteristics of the system itself. It is apparent that inventories are only one of the inputs and that the quantities needed depend importantly on the time lags -- procurement lead times, transportation times, transaction-processing times, repair-cycle times -- and on accounting accuracy and frequency of review.* A comprehensive analysis of spares management must consider all these inputs as substitutes for one another.

Inventory theory is also essential to this type of study, of course, because the comparison of alternative systems is not meaningful unless inventory levels are adapted to the characteristics of each. But studies oriented toward the design or modification of spares systems typically bring into focus a different type of problem and require different methods. Complete optimization of all the inputs or processes of such complex systems is not feasible, and the process of developing improved support structures is much more an engineering or design problem than a model-building task.

---

*These considerations are particularly important in view of the current rapid changes in computer technology and the increasing application of large-frame computers to inventory management.

It is sometimes fruitful to specify or design a number of alternative systems with differing characteristics and simulate the performance of each under the same conditions of demand or stress [22]. This procedure allows a much more elaborate representation of costs, priority rules, and interactions between system elements than any purely analytic approach. It may, for example, be important to study the effect of repair-capacity constraints, in conjunction with rules governing which spares items to repair first, on the performance of a multi-item stockage system. Stockage rules can be based on a relatively simple model of the process and tested in the more elaborate computer simulation, and alternative repair policies or capacities can be compared on the basis of cost and effectiveness.

A related kind of system analysis concerns the location of parts repair within a multi-point stockage system. Accomplishing repair at the lowest echelon reduces the stock-level requirements at those points by shortening the replenishment time but requires specialized equipment and other inputs to create the repair capability. The cost of alternative repair postures can be compared, providing that stock levels are adjusted to hold the effectiveness of the combined repair and stockage system constant for all comparisons.

Computation of Kits and War-reserve Levels

Another sort of stockage problem is represented by kits or reserve stocks for specific purposes; for example, the fly-away kit or mobility kit that is designed to provide spares support for some fixed period or number of missions, without resupply. Other related examples are the prestockage of war reserves at aircraft-dispersal sites and determination of shipboard-allowance kit quantities. In such problems the constraint generally is not cost, or at least not cost alone, but may instead be weight or cubage.

A particular problem connected with such kits is that, in most applications, the conditions under which spares-demand data are obtained can be expected to differ from those in which the kit will be used. One partial solution to this difficulty may be to collect

demand data from realistic exercises, but it is generally difficult to obtain large enough samples in this way. It is frequently necessary to use consumption data from routine operations, making judgment adjustments to reflect the likely conditions under which the kit is to be used. In this way, allowance can be made for the effect on spares demand of differing maintenance capabilities and, to some extent, for the effect of specific types of missions.

The inventory problem here typically takes the form of minimizing the expected shortages (or, in some versions, maximizing the probability of no shortage) for the value of the constraint. If J_i is a criticality factor and W_i the weight for the i^{th} part, the problem may take the form of minimizing the following for n candidates for the kit:*

$$(19) \qquad \sum_{i=1}^{n} J_i \sum_{x=S}^{\infty} (x-S)p_i(x;t)$$

Subject to

$$(20) \qquad \sum_{i=1}^{n} W_i S_i = W$$

This requires a computer solution analogous to that described for the allocation of investment in the multi-item (S-1,S) problem. Values of a constant, λ, are chosen and solutions found for the following i equations.

$$(21) \qquad \frac{\lambda W_i}{J_i} = \sum_{S_i^*}^{\infty} p_i(x;t)$$

Then $\sum_{i=1}^{n} W_i S_i$ is checked to see if it is greater or less than the constraint W, λ is revised accordingly, and (21) is resolved. It is usually possible to reach a satisfactory solution in a few iterations. As in previous examples, the interpretation of (21) must recognize that S_i is an integer. Linear programming methods can be used to optimize a

---

*The "kit problem" has received a number of formulations; this onc follows [16].

spares kit for more than one constraint; that is, a kit can be designed to minimize shortages for a target weight and volume or for weight, volume, and cost.

Determination of the essentiality or criticality weights is a large subject in itself. In some problems it is reasonable to assume that all candidates for a kit are equally essential, but in other applications differences in essentiality have been considered [5], [20]. Determining the weights involves the application of maintenance and operational expert judgment to derive a ranking of the candidate parts with regard to essentiality, considering the importance of each component to the mission, the degree to which it can be compensated for, if lost, by repair or redundant systems, and the urgency with which it must be replaced if defective. The parts are then grouped by major essentiality category and the weights, J_i, determined by judgment. The resulting kit can thus reflect, to some degree, the opinions of the users as to the relative importance of different spare parts to the mission for which the kit is designed.

Concluding Remarks

In this survey a great deal of attention has been given to theories of multi-item management, because such systems represent the most promising and useful tools currently being considered or applied. The major decisions of spares management -- the determination of the total requirement and the procurement budget -- pertain to the system as a whole, even though the budget must be implemented at the level of line-item detail. Management's concern is to find an allocation of its budget to spares purchases that will be reasonably efficient, that is, one under which no reallocation of the same budget could raise the overall effectiveness of the system. The ability to move between the domain of the individual spare part and that of an efficient system of many, related spare items is, perhaps, the main practical contribution that inventory studies can make at present.

But each approach to inventory control has its characteristic shortcomings as well as its advantages. As we have seen, it is possible to develop rather elaborate dynamic models of single items, which explicitly allow for varying programs and finite time horizons and embrace multiple points and echelon relationships. Thus, these include elements of reality not present in the multi-item techniques already discussed, useful though these techniques are. The search for ways of combining these two approaches -- the multi-item model for a single point under steady-state conditions and the complex dynamic model for a single item -- is an important research task. Bear in mind that computational practicality is always a required condition of a proposed solution for an inventory problem.

We have touched on another sense in which inventory analyses are limited representations of reality, without, perhaps, fully evaluating the consequences. This is the general area of prediction of demand and the related problem of forecasting the incidence of obsolescence. Spares-demand rates are often too low for satisfactory statistical prediction, and it seems quite likely that the underlying conditions governing demand are unstable for any very long forecast period.

Much progress has been made in this area, particularly in the development of Bayesian methods and the better understanding of multi-item models with aggregative criteria, but the demand problem remains as the least tractable (and most treacherous) aspect of inventory analyses generally. This is particularly true of the initial provisioning of spares, based on engineering estimates. Not only are the estimates as such highly subject to error, but design changes, with consequent obsolescence, are most common at this time. This suggests that an important characteristic of any inventory policy and the wider support framework in which it is embedded is its vulnerability to demand-prediction errors (or, more generally, to an improper specification of the demand model). Uncertainty is a dominating characteristic of the spares-management problem.

Such investigations will often lead to design studies of the whole support system, of the type described in the preceding section. In this connection, better methods are needed for studying large spares-support systems to find the preferred mix between inventory investment, data processing and management, transportation, and repair inputs. Ideally, "trade-off" studies of particular parts of the system should give way to a more general optimizing approach.

As a final comment, the criterion or performance measure for a spares inventory system is an arguable question. Judgmental considerations will influence the choice of both the criterion function itself and of the desired value of the function to be specified as part of the policy. This is because of the limited state of our knowledge of how spares availability, other logistics inputs, and the resultant military capability are related to each other. Simple aggregative measures such as the average number of units backordered or the fill rate appear to be fruitful criteria for suboptimizing the spare parts input, but more elaborate measures may eventually prove to be useful.

REFERENCES

1. Arrow, K. J., S. Karlin, and H. Scarf, Studies in the Mathematical Theory of Inventory and Production, Stanford University Press, Stanford, California, 1958.

2. Brown, R. G., Smoothing, Forecasting and Prediction of Discrete Time Series, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1962.

3. Clark, A. J., and H. Scarf, Optimal Policies for a Multi-Echelon Inventory Problem, PRC R-113, Planning Research Corporation, Los Angeles, 1959.

4. Davis, R. H., "Optimal Inventory Control Decision Rules for a Large Supply System," Operations Research, Vol. VII, No. 6, November-December 1959.

5. Denicoff, M., J. Fennell, S. E. Haber, W. H. Marlow, F. W. Segel, and Henry Solomon, The Polaris Military Essentiality System, T-171, George Washington University (Logistics Research Project), 1964.

6. Enke, S., "An Economist Looks at Air Force Logistics," The Review of Economics and Statistics, Vol. XL, No. 3, August 1958.

7. Enthoven, A. C., The Simple Mathematics of Maximization, P-1833, The RAND Corporation, Santa Monica, California, 1959.

8. Feeney, G. J., "A Basis for Strategic Decisions in Inventory Control Operations," Management Science, Vol. II, No. 1, October 1955.

9. -----, and C. C. Sherbrooke, The (s-1,s) Inventory Policy Under Compound Poisson Demand: A Theory of Recoverable Item Stockage, RM-4176-1-PR, The RAND Corporation, Santa Monica, California, 1966.

10. -----, An Objective Bayes Approach for Inventory Decisions, RM-4362-PR, The RAND Corporation, Santa Monica, California, 1965.

11. -----, A System Approach to Base Stockage of Recoverable Items, RM-4720-PR, The RAND Corporation, Santa Monica, California, 1965.

12. Feller, W., An Introduction to Probability Theory and Its Applications, Vol. I (2nd ed.), John Wiley & Sons, Inc., New York, 1961.

13. Ferguson, A. R., and L. Fisher, Stockage Policy for Medium and Low-Cost Parts, RM-1962, The RAND Corporation, Santa Monica, California, 1958.

14. Fukuda, Y., Bayes and Maximum Likelihood Policies for a Multi-Echelon Inventory Problem, PRC R-161, Planning Research Corporation, Los Angeles, California, 1960.

15. Galliher, H. P., P. M. Morse, and M. Simond, "Dynamics of Two Classes of Continuous-Review Inventory Systems," Operations Research, Vol. VII, No. 3, May-June 1959.

16. Geisler, M. A., and H. W. Karr, "The Design of Military Supply Tables for Spare Parts," Operations Research, Vol. IV, No. 4, August 1956.

17. Gradwohl, A. J., Case Studies in the Multi-Echelon Inventory Problem, PRC R-133, Planning Research Corporation, Los Angeles, California, 1959.

18. Hitch, C. J., "Suboptimization in Operations Problems," Operations Research, Vol. I, No. 4, August 1953.

19. McGlothlin, W. H., Development of Bayesian Parameters for Spare-Parts Demand Prediction, RM-3699-PR, The RAND Corporation, Santa Monica, California, 1963.

20. Okun, B., "Design, Test, and Evaluation of an Experimental Flyaway Kit," Naval Research Logistics Quarterly, Vol. VII, No. 2, June 1960.

21. Petersen, J. W., Savings from Procurement Deferral with Interim Contractor Support: The Case of High Value Airframe Spares, RM-2085, The RAND Corporation, Santa Monica, California, 1958.

22. -----, and H. W. Nelson, Integrated Materiel Management, RM-2870-PR, The RAND Corporation, Santa Monica, California, 1961.

23. Raiffa, H., and R. Schlaifer, Applied Statistical Decision Theory, Division of Research, Harvard Business School, 1961.

24. Scarf, H., "The Optimality of (S,s) Policies in the Dynamic Inventory Problem," Mathematical Methods in the Social Sciences, 1959, (K. J. Arrow, S. Karlin, and P. Suppes, eds.), Stanford University Press, Stanford, California, 1960.

25. Sherbrooke, C. C., METRIC: A Multi-Echelon Technique for Recoverable Item Control, RM-5078-PR, The RAND Corporation, 1966.

26. Whitin, T. M., and G. Hadley, Analysis of Inventory Systems, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1963.